

Claims

What is claimed is:

1. A method for performing an outer join of at least a first table T1 and a second table T2, where  
5 the join has join conditions, each of the tables has an associated Star Map, S1 and S2, respectively, and each Star Map includes bitmap entries having locations indexed by the hash of one or more values associated with one or more join key columns of its associated table, where a bitmap entry in a Star Map, if set, indicates the presence of a row in its associated table that has entries in the one or more join key columns that together hash to the location of the bitmap entry, the method including:

- a) performing one or more Boolean operations using the bitmap entries of the Star Maps S1 and S2 to produce set bitmap entries in a Star Map SJ where there is a corresponding set bitmap entry in S1 and a corresponding set bitmap entry in S2;
- b) selecting a row from table T1 and hashing the combined entries in the one or more join key columns of the selected T1 row to identify a bitmap entry in SJ;
- c) if the identified bitmap entry in SJ is not set, projecting the selected T1 row with a NULL corresponding to data from table T2; and
- d) if the identified bitmap entry in SJ is set, performing the following:
  - d1) if no row in T2 satisfies the join conditions and has entries in its one or more join key columns that together hash to the location of the identified set bitmap entry in SJ, projecting the selected T1 row and a NULL corresponding to data from table T2; and
  - d2) otherwise, for each row from T2 that satisfies the join condition and has entries in its one or more join key columns that together hash to the location of the identified set bitmap entry in SJ, projecting the selected T1 row with data from the row from T2;
- e) repeating b)-d) for all rows in T1.

2. The method of claim 1, wherein a plurality of tables T2' and a plurality of associated Star Maps S2' are provided, whereby Star Map S1 is logically ANDed with each Star Map S2' to generate join Star Maps SJ', respectively, and (b) through (e) are executed for all tables T2' and associated Star Maps S2'.

5 3. The method of claim 1 further comprising determining the expected cardinality of the join result, and if the cardinality is less than a predefined threshold value, performing a) through e).

4. The method of claim 3, wherein the threshold value is determined dynamically depending on at least one parameter.

5. The method of claim 1, wherein T1 is the right table in a right outer join operation.

6. The method of claim 1, wherein T1 is the left table in a left outer join operation.

7. The method of claim 1, wherein the one or more Boolean operations is a logical AND operation.

8. A computer program, stored on a tangible storage medium, for performing an outer join of at least a first table T1 and a second table T2, where the join has join conditions, each of the tables has an associated Star Map, S1 and S2, respectively, and each Star Map includes bitmap entries having locations indexed by the hash of one or more values associated with one or more join key columns of its associated table, where a bitmap entry in a Star Map, if set, indicates the presence of a row in its associated table that has entries in the one or more join key columns that together hash to the location of the bitmap entry, the program including executable instructions that cause a computer to:

5       a) perform one or more Boolean operations using the bitmap entries of the Star Maps S1 and S2 to produce set bitmap entries in a Star Map SJ where there is a corresponding set bitmap entry in S1 and a corresponding set bitmap entry in S2;

10      b) select a row from table T1 and hash the combined entries in the one or more join key columns of the selected T1 row to identify a bitmap entry in SJ;

15      c) if the identified bitmap entry in SJ is not set, project the selected T1 row with a NULL corresponding to data from table T2; and

20      d) if the identified bitmap entry in SJ is set, performing the following:

25        d1) if no row in T2 satisfies the join conditions and has entries in its one or more join key columns that together hash to the location of the identified set bitmap entry in SJ, project the selected T1 row and a NULL corresponding to data from table T2; and

30        d2) otherwise, for each row from T2 that satisfies the join conditions and has entries in its one or more join key columns that together hash to the location of the identified set bitmap entry in SJ, project the selected T1 row with data from the row from T2;

25      e) repeat b)-d) for all rows in T1.

9. The computer program of claim 8, wherein a plurality of tables T2' and a plurality of associated Star Maps S2' are provided, and the program includes executable instructions for causing the Star Map S1 to be logically ANDed with each Star Map S2' to generate join Star Maps SJ', respectively, and the program executes (b) through (e) for all tables T2' and associated 30 Star Maps S2'.

10. The computer program of claim 8 wherein the program includes executable instructions for determining the cardinality of the join result, and if the cardinality is less than a predefined threshold value, performing a) through e).

11. The computer program of claim 10, wherein the threshold value is determined dynamically depending on at least one parameter.

5

12. The method of claim 8, wherein T1 is the right table in a right outer join operation.

13. The method of claim 8, wherein T1 is the left table in a left outer join operation.

14. The method of claim 8, wherein the one or more Boolean operations is a logical AND operation.

10237-001

15. A database system for accessing a database according to a outer join query, the query including join conditions, the database system including

5      a massively parallel processing system including

10      one or more nodes;

15      a plurality of CPUs, each of the one or more nodes providing access to one or more CPUs;

20      a plurality of virtual processes each of the one or more CPUs providing access to one or more processes;

25      each process configured to manage data stored in one of a plurality of data-storage facilities;

30      at least a first table T1 and a second table T2 being distributed among the data-storage facilities;

35      each of the tables having an associated Star Map, S1 and S2, respectively, each Star Map being distributed among the data-storage facilities,

40      each Star Map including bitmap entries having locations indexed by the hash of one or more values associated with one or more join key columns of its associated table, where a bitmap entry in a Star Map, if set, indicates the presence of a row in its associated table that has entries in the one or more join key columns that together hash to the location of the bitmap entry;

45      a join process executed on one or more of the plurality of CPUs that cause the CPUs to

50      a) perform one or more Boolean operations using the bitmap entries of the Star Maps S1 and S2 to produce set bitmap entries in a Star Map SJ where there is a corresponding set bitmap entry in S1 and a corresponding set bitmap entry in S2;

55      b) select a row from table T1 and hash the combined entries in the one or more join key columns of the selected T1 row to identify a bitmap entry in SJ;

60      c) if the identified bitmap entry in SJ is not set, project the selected T1 row with a NULL corresponding to data from table T2; and

65      d) if the identified bitmap entry in SJ is set, performing the following:

10  
15  
20

d1) if no row in T2 satisfies the join conditions and has entries in its one or more join key columns that together hash to the location of the identified set bitmap entry in SJ , project the selected T1 row and a NULL corresponding to data from table T2; and

5 d2) otherwise, for each row from T2 that satisfies the join condition and has entries in its one or more join key columns that together hash to the location of the identified set bitmap entry in SJ, project the selected T1 row with data from the row from T2,

e) repeat b)-d) for all rows in T1.

16. The database system of claim 15, wherein a plurality of tables T2' and a plurality of associated Star Maps S2' are provided, and the join process further causes the CPUs to logically AND the Star Map S1 with each Star Map S2' to generate join Star Maps SJ', respectively, and to execute (b) through (e) for all tables T2' and associated Star Maps S2'.

17. The database system of claim 15 wherein the join process further causes the CPUs to determine the cardinality of the join result, and if the cardinality is less than a predefined threshold value, performing a) through e).

18. The database system of claim 17, wherein the join process further causes the CPUs to determine the threshold value dynamically depending on at least one parameter.

19. The database system of claim 15, wherein T1 is the right table in a right outer join operation.

20. The database system of claim 15, wherein T1 is the left table in a left outer join operation.

21. The database system of claim 15, wherein the one or more Boolean operations is a logical AND operation.